

Optimism in Reinforcement Learning and Kullback-Leibler Divergence

Sarah Filippi, Olivier Cappé and Aurélien Garivier
LTCI, TELECOM ParisTech and CNRS
46 rue Barrault, 75013 Paris, France
Email: (filippi, cappe, garivier)@telecom-paristech.fr*

September 25, 2011

Abstract

We consider model-based reinforcement learning in finite Markov Decision Processes (MDPs), focussing on so-called optimistic strategies. In MDPs, optimism can be implemented by carrying out extended value iterations under a constraint of consistency with the estimated model transition probabilities. The UCRL2 algorithm by Auer, Jaksch and Ortner (2009), which follows this strategy, has recently been shown to guarantee near-optimal regret bounds. In this paper, we strongly argue in favor of using the Kullback-Leibler (KL) divergence for this purpose. By studying the linear maximization problem under KL constraints, we provide an efficient algorithm, termed KL-UCRL, for solving KL-optimistic extended value iteration. Using recent deviation bounds on the KL divergence, we prove that KL-UCRL provides the same guarantees as UCRL2 in terms of regret. However, numerical experiments on classical benchmarks show a significantly improved behavior, particularly when the MDP has reduced connectivity. To support this observation, we provide elements of comparison between the two algorithms based on geometric considerations. Keywords : Reinforcement learning; Markov decision processes; Model-based approaches; Optimism; Kullback-Leibler divergence; Regret bounds

1 Introduction

In reinforcement learning, an agent interacts with an unknown environment, aiming to maximize its long-term payoff [15]. This interaction is modelled by a Markov Decision Process (MDP) and it is assumed that the agent does not know the parameters of the process and needs to learn directly from observations. The agent thus faces a fundamental trade-off between gathering experimental data about the consequences of the actions (exploration) and acting consistently with

*This work has been partially supported by Orange Labs under contract n°289365.

past experience in order to maximize the rewards (exploitation).

We consider in this article a MDP with finite state and action spaces for which we propose a *model-based* reinforcement learning algorithm, i.e., an algorithm that maintains running estimates of the model parameters (transitions probabilities and expected rewards) [6, 10, 14, 16]. A well-known approach to balance exploration and exploitation, followed for example by the well-know algorithm R-MAX [4], is the so-called *optimism in the face of uncertainty* principle. It was first proposed in the multi-armed bandit context by [11], and has been extended since then to several frameworks: instead of acting optimally according to the estimated model, the agent follows the optimal policy for a surrogate model, named *optimistic model*, which is close enough to the former but leads to a higher long-term reward. The performance of such an algorithm can be analyzed in terms of *regret*, which consists in comparing the rewards collected by the algorithm with the rewards obtained when following an optimal policy. The study of the asymptotic regret due to [11] in the multi-armed context has been extended to MDPs by [5], proving that an optimistic algorithm can achieve logarithmic regret. The subsequent works [2, 9, 3] introduced algorithms that guarantee non-asymptotic logarithmic regret in a large class of MDPs. In these latter works, the optimistic model is computed using the L^1 (or total variation) norm as a measure of proximity between the estimated and optimistic transition probabilities.

In addition to logarithmic regret bounds, the UCRL2 algorithm of [9] is also attractive due to the simplicity of each L^1 extended value iteration step. In this case, optimism simply results in adding a bonus to the most promising transition (i.e., the transition that leads to the state with current highest value) while removing the corresponding probability mass from less promising transitions. This process is both elementary and easily interpretable, which is desirable in some applications.

However, the L^1 extended value iteration leads to undesirable pitfalls, which may compromise the practical performance of the algorithm. First, the optimistic model is not continuous with respect to the estimated parameters – small changes in the estimates may result in very different optimistic models. More importantly, the L^1 optimistic model can become incompatible with the observations by assigning a probability of zero to a transition that has actually been observed. Moreover, in MDPs with reduced connectivity, L^1 optimism results in a persistent bonus for all transitions heading towards the most valuable state, even when significant evidence has been accumulated that these transitions are impossible.

In this paper, we propose an improved optimistic algorithm, called KL-UCRL, that avoids these pitfalls altogether. The key is the use of the Kullback-Leibler (KL) pseudo-distance instead of the L^1 metric, as in [5]. Indeed, the smoothness of the KL metric largely alleviates the first issue. The second issue is completely avoided thanks to the strong relationship between the geometry of the probability simplex induced by the KL pseudo-metric and the theory of large deviations. For the third issue, we show that the KL-optimistic model results from a trade-off between the relative value of the most promising state

and the statistical evidence accumulated so far regarding its reachability.

We provide an efficient procedure, based on one-dimensional line searches, to solve the linear maximization problem under KL constraints. As a consequence, the numerical complexity of the KL-UCRL algorithm is comparable to that of UCRL2. Building on the analysis of [9, 3, 1], we also obtain logarithmic regret bounds for the KL-UCRL algorithm. The proof of this result is based on novel concentration inequalities for the KL-divergence, which have interesting properties when compared with those traditionally used for the L^1 norm. Although the obtained regret bounds are comparable to earlier results in term of rate and dependence in the number of states and actions, we observed in practice significant performance improvements. This observation is illustrated using benchmark examples (the *RiverSwim* and *SixArms* environments of [14]) and through a thorough discussion of the geometric properties of KL neighborhoods.

The paper is organized as follows. The model and a brief survey of the value iteration algorithm in undiscounted MDPs are presented in Section 2. Section 3 and 4 are devoted, respectively, to the description and the analysis of the KL-UCRL algorithm. Section 5 contains numerical experiments and Section 6 concludes the paper by discussing the advantages of using KL rather than L^1 confidence neighborhoods.

2 Markov Decision Process

Consider a Markov decision process (MDP) $\mathbf{M} = (\mathcal{X}, \mathcal{A}, P, r)$ with finite state space \mathcal{X} , and action space \mathcal{A} . Let $X_t \in \mathcal{X}$ and $A_t \in \mathcal{A}$ denote respectively the state of the system and the action chosen by the agent at time t . The probability to jump from state X_t to state X_{t+1} is denoted by $P(X_{t+1}; X_t, A_t)$. Besides, the agent receives at time t a random reward $R_t \in [0, 1]$ with mean $r(X_t, A_t)$. The aim of the agent is to choose the sequence of actions so as to maximize the cumulated reward. His choices are summarized in a *stationary policy* $\pi : \mathcal{X} \rightarrow \mathcal{A}$.

In this paper, we consider *communicating* MDPs, i.e., MDPs such that for any pair of states x, x' , there exists policies under which x' can be reached from x with positive probability. For those MDPs, it is known that the *average reward* following a stationary policy π , denoted by $\rho^\pi(\mathbf{M})$ and defined as

$$\rho^\pi(\mathbf{M}) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}_{\mathbf{M}, \pi} \left(\sum_{t=0}^n R_t \right),$$

is state-independent [13]. Let $\pi^*(\mathbf{M}) : \mathcal{X} \rightarrow \mathcal{A}$ and $\rho^*(\mathbf{M})$ denote respectively the optimal policy and the optimal average reward: $\rho^*(\mathbf{M}) = \sup_{\pi} \rho^\pi(\mathbf{M}) = \rho^{\pi^*(\mathbf{M})}(\mathbf{M})$. The notations $\rho^*(\mathbf{M})$ and $\pi^*(\mathbf{M})$ are meant to highlight the fact that both the optimal average reward and the optimal policy depend on the model \mathbf{M} . The optimal average reward satisfies the so-called *Bellman optimality*

equation: for all $x \in \mathcal{X}$,

$$h^*(\mathbf{M}, x) + \rho^*(\mathbf{M}) = \max_{a \in \mathcal{A}} \left(r(x, a) + \sum_{x' \in \mathcal{X}} P(x'; x, a) h^*(\mathbf{M}, x') \right),$$

where the $|\mathcal{X}|$ -dimensional vector $h^*(\mathbf{M})$ is called a *bias* vector. Note that it is only defined up to an additive constant. For a fixed MDP \mathbf{M} , the optimal policy $\pi^*(\mathbf{M})$ can be derived by solving the optimality equation and by defining, for all $x \in \mathcal{X}$,

$$\pi^*(\mathbf{M}, x) \in \operatorname{argmax}_{a \in \mathcal{A}} \left(r(x, a) + \sum_{x' \in \mathcal{X}} P(x'; x, a) h^*(\mathbf{M}, x') \right).$$

In practice, the optimal average reward and the optimal policy may be computed, for instance, using the value iteration algorithm [13].

3 The KL-UCRL algorithm

In this paper, we focus on the reinforcement learning problem in which the agent does not know the model \mathbf{M} beforehand, i.e. the transition probabilities and the distribution of the rewards are unknown. More specifically, we consider model-based reinforcement learning algorithms which estimate the model through observations and act accordingly. Denote by $\hat{P}_t(x'; x, a)$ the estimate at time t of the transition probability from state x to state x' conditionally to the action a , and, by $\hat{r}_t(x, a)$ the mean reward received in state x when action a has been chosen. We have:

$$\begin{aligned} \hat{P}_t(x'; x, a) &= \frac{N_t(x, a, x')}{\max(N_t(x, a), 1)} \\ \hat{r}_t(x, a) &= \frac{\sum_{k=0}^{t-1} R_k \mathbf{1}_{\{X_k=x, A_k=a\}}}{\max(N_t(x, a), 1)}, \end{aligned} \quad (1)$$

where $N_t(x, a, x') = \sum_{k=0}^{t-1} \mathbf{1}_{\{X_k=x, A_k=a, X_{k+1}=x'\}}$ is the number of visits, up to time t , to the state x followed by a visit to x' when the action a has been chosen, and similarly, $N_t(x, a) = \sum_{k=0}^{t-1} \mathbf{1}_{\{X_k=x, A_k=a\}}$. The optimal policy in the estimated model $\widehat{\mathbf{M}}_t = (\mathcal{X}, \mathcal{A}, \hat{P}_t, \hat{r}_t)$ may be misleading due to estimation errors: pure exploitation policies are commonly known to fail with positive probability. To avoid this problem, *optimistic model-based approaches* consider a set \mathcal{M}_t of potential MDPs including $\widehat{\mathbf{M}}_t$ and choose the MDP from this set that leads to the largest average reward. In the following, the set \mathcal{M}_t is defined as follows:

$$\begin{aligned} \mathcal{M}_t &= \{ \mathbf{M} = (\mathcal{X}, \mathcal{A}, P, r) : \forall x \in \mathcal{X}, \forall a \in \mathcal{A}, \\ &\quad |\hat{r}_t(x, a) - r(x, a)| \leq \epsilon_R(x, a, t) \\ &\quad \text{and } d(\hat{P}_t(\cdot; x, a), P(\cdot; x, a)) \leq \epsilon_P(x, a, t) \}, \end{aligned}$$

where d measures the difference between the transition probabilities. The *radius of the neighborhoods* $\epsilon_R(x, a, t)$ and $\epsilon_P(x, a, t)$ around, respectively, the estimated reward $\hat{r}_t(x, a)$ and the estimated transition probabilities $\hat{P}_t(\cdot; x, a)$, decrease with $N_t(x, a)$.

In contrast to UCRL2, which uses the L^1 -distance for d , we propose to rely on the Kullback-Leibler divergence, as in the seminal article [5]; however, contrary to the approach of [5], no prior knowledge on the state structure of the MDP is needed. Recall that the Kullback-Leibler divergence is defined for all n -dimensional probability vectors p and q by $KL(p, q) = \sum_{i=1}^n p_i \log \frac{p_i}{q_i}$ (with the convention that $0 \log 0 = 0$). In the sequel, we will show that this choice dramatically alters the behavior of the algorithm and leads to significantly better performance, while causing a limited increase of complexity; in Section 6, the advantages of using a KL-divergence instead of the L^1 -norm are illustrated and argued.

3.1 The KL-UCRL algorithm

The KL-UCRL algorithm, described below, is a variant of the efficient model-based algorithm UCRL2, introduced by [1] and extended to more general MDPs by [3]. The key step of the algorithm, the search for the optimistic model (Step 8), is detailed below as Algorithm 2.

Algorithm 1 KL-UCRL

- 1: Initialization: $j = 0$, $t_0 = 0$; $\forall a \in \mathcal{A}, \forall x \in \mathcal{X}, n_0(x, a) = 0$, $N_0(x, a) = 0$; initial policy π_0 .
- 2: **for** all $t \geq 1$ **do**
- 3: Observe X_t
- 4: **if** $n_j(X_t, \pi_j(X_t)) \geq \max(N_{t_j}(X_t, \pi_j(X_t)), 1)$ **then**
- 5: *Begin a new episode:* $j = j + 1$, $t_j = t$,
- 6: Reinitialize: $\forall a \in \mathcal{A}, \forall x \in \mathcal{X}$, $n_j(x, a) = 0$
- 7: Estimate \hat{P}_t and \hat{r}_t according to (1)
- 8: Find the optimistic model $\mathbf{M}_j \in \mathcal{M}_t$ and the related policy π_j solving equation (2) and using Algorithm 2
- 9: **end if**
- 10: Choose action $A_t = \pi_j(X_t)$
- 11: Receive reward R_t
- 12: Update the count within the current episode:

$$n_j(X_t, A_t) = n_j(X_t, A_t) + 1$$

- 13: Update the global count:

$$N_t(X_t, A_t) = N_{t-1}(X_t, A_t) + 1$$

- 14: **end for**
-

The KL-UCRL algorithm proceeds in episodes. Let t_j be the starting time of episode j ; the length of the j -th episode depends on the number of visits $N_{t_j}(x, a)$ to each state-action pair (x, a) before t_j compared to the number of visits $n_j(x, a)$ to the same pair during the j -th episode. More precisely, an episode ends as soon as $n_j(x, a) \geq N_{t_j}(x, a)$ for some state-action pair (x, a) . The policy π_j , followed during the j -th episode, is an optimal policy for the optimistic MDP $\mathbf{M}_j = (\mathcal{X}, \mathcal{A}, P_j, r_j) \in \mathcal{M}_{t_j}$, which is computed by solving the *extended optimality equations*: for all $x \in \mathcal{X}$

$$h^*(x) + \rho^* = \max_{P, r} \max_{a \in \mathcal{A}} \left(r(x, a) + \sum_{x' \in \mathcal{X}} P(x'; x, a) h^*(x') \right) \quad (2)$$

where the maximum is taken over all P, r such that

$$\begin{aligned} \forall x, \forall a, \quad KL(\hat{P}_{t_j}(\cdot; x, a), P(\cdot; x, a)) &\leq \frac{C_P}{N_{t_j}(x, a)}, \\ \forall x, \forall a, \quad |\hat{r}_{t_j}(x, a) - r(x, a)| &\leq \frac{C_R}{\sqrt{N_{t_j}(x, a)}}, \end{aligned}$$

where C_P and C_R are constants which control the size of the confidence balls. The transition matrix P_j and the mean reward r_j of the optimistic MDP \mathbf{M}_j maximize those equations. The *extended value iteration* algorithm may be used to approximately solve the fixed point equation (2) [13, 1].

3.2 Maximization of a linear function on a KL-ball

At each step of the extended value iteration algorithm, the maximization problem (2) has to be solved. For every state x and action a , the maximization of $r(x, a)$ under the constraint that $|\hat{r}_{t_j}(x, a) - r(x, a)| \leq C_R/\sqrt{N_{t_j}(x, a)}$ is obviously solved taking $r(x, a) = \hat{r}_{t_j}(x, a) + C_R/\sqrt{N_{t_j}(x, a)}$, so that the main difficulty lies in maximizing the dot product between the probability vector $q = P(\cdot; x, a)$ and the *value vector* $V = h^*$ over a KL-ball around the fixed probability vector $p = \hat{P}_{t_j}(\cdot; x, a)$:

$$\max_{q \in \mathbb{S}^{|\mathcal{X}|}} V'q \quad \text{s.t.} \quad KL(p, q) \leq \epsilon, \quad (3)$$

where V' denotes the transpose of V and \mathbb{S}^n the set of n -dimensional probability vectors. The radius of the neighborhood $\epsilon = C_P/N_{t_j}(x, a)$ controls the size of the confidence ball. This convex maximization problem is studied in Appendix A, leading to the efficient algorithm presented below. Detailed analysis of the Lagrangian of (3) shows that the solution of the maximization problem essentially relies on finding roots of the function f (that depends on the parameter V), defined as follows: for all $\nu \geq \max_{i \in \bar{Z}} V_i$, with $\bar{Z} = \{i : p_i > 0\}$,

$$f(\nu) = \sum_{i \in \bar{Z}} p_i \log(\nu - V_i) + \log \left(\sum_{i \in \bar{Z}} \frac{p_i}{\nu - V_i} \right). \quad (4)$$

In the special case where the most promising state i_M has never been reached from the current state-action pair (i.e. $p_{i_M} = 0$), the algorithm makes a trade-off between the relative value of the most promising state V_{i_M} and the statistical evidence accumulated so far regarding its reachability.

Algorithm 2 Function MaxKL

Input A value function V , a probability vector p , a constant ϵ

Output A probability vector q that maximizes (3)

- 1: Let $Z = \{i : p_i = 0\}$ and $\bar{Z} = \{i : p_i > 0\}$.
- 2: Let $I^* = Z \cap \operatorname{argmax}_i V_i$
- 3: **if** $I^* \neq \emptyset$ and there exists $i \in I^*$ such that $f(V_i) < \epsilon$ **then**
- 4: Let $\nu = V_i$ and $r = 1 - \exp(f(\nu) - \epsilon)$.
- 5: For all $i \in I^*$, assign values of q_i such that

$$\sum_{i \in I^*} q_i = r .$$

- 6: For all $i \in Z/I^*$, let $q_i = 0$.
 - 7: **else**
 - 8: For all $i \in Z$, let $q_i = 0$. Let $r = 0$.
 - 9: Find ν such that $f(\nu) = \epsilon$ using Newton's method.
 - 10: **end if**
 - 11: For all $i \in \bar{Z}$, let $q_i = \frac{(1-r)\tilde{q}_i}{\sum_{i \in \bar{Z}} \tilde{q}_i}$ where $\tilde{q}_i = \frac{p_i}{\nu - V_i}$.
-

In practice, f being a convex positive decreasing function (see Appendix B), Newton's method can be applied to find ν such that $f(\nu) = \epsilon$ (in Step 9 of the algorithm), so that numerically solving (3) is a matter of a few iterations. Appendix B contains a discussion of the initialization of Newton's algorithm based on asymptotic arguments.

4 Regret bounds

4.1 Theoretical results

To analyze the performance of KL-UCRL, we compare the rewards accumulated by the algorithm to the rewards that would be obtained, on average, by an agent playing an optimal policy. The *regret* of the algorithm after T steps is defined as in [9]:

$$\text{Regret}_T = \sum_{t=1}^T (\rho^*(\mathbf{M}) - R_t) .$$

We adapt the regret bound analysis of the UCRL2 algorithm to the use of KL-neighborhoods, and obtain similar theorems. Let

$$D(\mathbf{M}) = \max_{x, x'} \min_{\pi} \mathbb{E}_{\mathbf{M}, \pi}(\tau(x, x')) ,$$

where $\tau(x, x')$ is the hitting time of x' , starting from state x . The $D(\mathbf{M})$ constant will appear in the regret bounds. For all communicating MDPs \mathbf{M} , $D(\mathbf{M})$ is finite. Theorem 1 establishes an upper bound on the regret of the KL-UCRL algorithm with C_P and C_R defined as

$$C_P = |\mathcal{X}| \left(B + \log \left(B + \frac{1}{\log(T)} \right) \left[1 + \frac{1}{B + \frac{1}{\log(T)}} \right] \right)$$

where $B = \log \left(\frac{2e|\mathcal{X}|^2|\mathcal{A}|\log(T)}{\delta} \right)$ and

$$C_R = \sqrt{\frac{\log(4|\mathcal{X}||\mathcal{A}|\log(T)/\delta)}{1.99}}.$$

Theorem 1 *With probability $1 - \delta$, it holds that for $T > 5$, the regret of KL-UCRL is bounded by*

$$\text{Regret}_T \leq CD(\mathbf{M})|\mathcal{X}|\sqrt{|\mathcal{A}|T\log(\log(T)/\delta)},$$

for a constant $C \leq 24$ that does not depend on the model.

It is also possible to prove a logarithmic upper bound for the expected regret. This bound, presented in Theorem 2, depends on the model through the constant $\Delta(\mathbf{M})$ defined as $\Delta(\mathbf{M}) = \rho^*(\mathbf{M}) - \max_{\pi, \rho^\pi(\mathbf{M}) < \rho^*(\mathbf{M})} \rho^\pi(\mathbf{M})$. $\Delta(\mathbf{M})$ quantifies the margin between optimal and suboptimal policies.

Theorem 2 *For $T > 5$, the expected regret of KL-UCRL is bounded by*

$$\mathbb{E}(\text{Regret}_T) \leq CD^2(\mathbf{M}) \frac{|\mathcal{X}|^2|\mathcal{A}|\log(T)}{\Delta(\mathbf{M})} + C(\mathbf{M}),$$

where $C \leq 400$ is a constant independent of the model, and $C(\mathbf{M})$ is a constant which depends on the model (see [9]).

4.2 Elements of proof

The proof of Theorem 1 is inspired from [9, 3]. Due to the lack of space, we only provide the main steps of the proof. First, the following proposition enables us to ensure that, with high probability, the true model $\mathbf{M} = (\mathcal{X}, \mathcal{A}, P, r)$ belongs to the set of models \mathcal{M}_t at each time step.

Proposition 1 *For every horizon $T \geq 1$ and for $\delta > 0$, $\mathbb{P}(\forall t \leq T, \mathbf{M} \in \mathcal{M}_t) \geq 1 - 2\delta$.*

The proof relies on the two following concentration inequalities due to [7, 8]: for all $x \in \mathcal{X}$, $a \in \mathcal{A}$, any $C_P > 0$, and $C_R > 0$, it holds that

$$\begin{aligned} & \mathbb{P} \left(\forall t \leq T, \text{KL}(\hat{P}_t(\cdot; x, a), P(\cdot; x, a)) > \frac{C_P}{N_t(x, a)} \right) \\ & \leq 2e(C_P \log(T) + |\mathcal{X}|)e^{-\frac{C_P}{|\mathcal{X}|}} = 1 - \frac{\delta}{|\mathcal{X}||\mathcal{A}|} \\ & \mathbb{P} \left(\forall t \leq T, |\hat{r}_t(x, a) - r(x, a)| \leq \frac{C_R}{\sqrt{N_t(x, a)}} \right) \\ & \leq 4\log(T)e^{-1.99C_R} = 1 - \frac{\delta}{|\mathcal{X}||\mathcal{A}|}. \end{aligned} \quad (5)$$

Then, summing over all state-action pairs, Proposition 1 follows.

Using Hoeffding's inequality, with high probability, the regret at time T can be written as the sum of a regret in each of the $m(T)$ episodes plus an additional term $C_e(T, \delta) = \sqrt{T \log(1/\delta)/2}$:

$$\begin{aligned} \text{Regret}_T & \leq \sum_{(x,a)} N_T(x, a)(\rho^*(\mathbf{M}) - r(x, a)) + C_e(T, \delta) \\ & \leq \sum_{k=1}^{m(T)} \sum_{(x,a)} n_k(x, a)(\rho^*(\mathbf{M}) - r(x, a)) + C_e(T, \delta) \end{aligned}$$

Let P_k and π_k denote, respectively, the transition probability matrix of the optimistic model and the optimal policy in the k -th episode ($1 \leq k \leq m(T)$). It is easy to show that (see [9] for details), with probability $1 - \delta$,

$$\begin{aligned} \text{Regret}_T & \leq \sum_{k=1}^{m(T)} \sum_{x \in \mathcal{X}} n_k(x, \pi_k(x)) [\\ & \quad + (r_k(x, \pi_k(x)) - r(x, \pi_k(x))) \\ & \quad + (P_k(\cdot; x, \pi_k(x)) - P(\cdot; x, \pi_k(x)))' h_k \\ & \quad + (P(\cdot; x, \pi_k(x)) - \mathbf{e}_x)' h_k] \\ & \quad + C_e(T, \delta), \end{aligned}$$

where h_k is a bias vector, $\mathbf{e}_x(y) = 1$ if $x = y$ and $\mathbf{e}_x(y) = 0$ otherwise. We now bound each of the three terms in the previous summation. Denote by $n_k^{\pi_k}$ the row vector such that $n_k^{\pi_k}(x) = n_k(x, \pi_k(x))$ and by $r_k^{\pi_k}$ (resp. r^{π_k}) the column vector such that $r_k^{\pi_k}(x) = r_k(x, \pi_k(x))$ (resp. $r^{\pi_k}(x) = r(x, \pi_k(x))$). Similarly $P_k^{\pi_k}$ (resp. P^{π_k}) is the transition matrix if the policy π_k is followed under the optimistic model \mathbf{M}_k (resp. the true model \mathbf{M}). If the true model $\mathbf{M} \in \mathcal{M}_{t_k}$, we have for all $x \in \mathcal{X}$, for all $a \in \mathcal{A}$,

$$n_k^{\pi_k}(r_k^{\pi_k} - r^{\pi_k}) \leq 2 \sum_{(x,a)} n_k(x, a) \sqrt{\frac{C_R}{N_t(x, a)}} \quad (6)$$

Using Pinsker's inequality, and the fact that $\|h_k\|_\infty \leq D$ [9],

$$\begin{aligned}
& n_k^{\pi_k} (P_k^{\pi_k} - P^{\pi_k}) h_k \\
& \leq \sum_{(x,a)} n_k(x,a) \|P_k(\cdot; x, a) - P(\cdot; x, a)\|_1 \|h_k\|_\infty \\
& \leq 2D\sqrt{2} \sum_{(x,a)} n_k(x,a) \sqrt{\frac{C_P}{N_{t_k}(x,a)}}.
\end{aligned} \tag{7}$$

The third term $n_k^{\pi_k} (P^{\pi_k} - I) h_k$ may be written as follows:

$$\begin{aligned}
n_k^{\pi_k} (P^{\pi_k} - I) h_k & \leq \sum_{t=t_k}^{t_{k+1}-1} (P(\cdot; X_t, A_t) - \mathbf{e}_{X_{t+1}}) h_k \\
& \quad + h_k(X_{t+1}) - h_k(X_t),
\end{aligned}$$

where \mathbf{e}_x is the all 0's vector with a 1 only on the x -th component. For all $t \in [t_k, t_{k+1}-1]$, note that $\xi_t = (P(\cdot; X_t, A_t) - \mathbf{e}_{X_{t+1}}) h_k$ is a martingale difference upper-bounded by D . Applying the Azuma-Hoeffding inequality, we obtain that

$$\begin{aligned}
\sum_{k=1}^{m(T)} n_k^{\pi_k} (P^{\pi_k} - I) h_k & = \sum_{t=1}^T \xi_t + m(T)D \\
& \leq D\sqrt{\frac{T \log(1/\delta)}{2}} + m(T)D
\end{aligned} \tag{8}$$

with probability $1 - \delta$. In addition, Auer and al [1] proved that

$$\sum_{k=1}^{m(T)} \sum_{x,a} \frac{n_k(x,a)}{\sqrt{N_{t_k}(x,a)}} \leq (\sqrt{2} + 1) \sqrt{|\mathcal{X}||\mathcal{A}|T}$$

and

$$m(T) \leq |\mathcal{X}||\mathcal{A}| \log_2 \left(\frac{8T}{|\mathcal{X}||\mathcal{A}|} \right).$$

Combining all the terms completes the proof of Theorem 1. The proof of Theorem 2 follows from Theorem 1 using the same arguments as in the proof of Theorem 4 in [9].

5 Numerical experiments

To compare the behavior of algorithms KL-UCRL and UCRL2, we consider the benchmark environments *RiverSwim* and *SixArms* proposed by [14] as well as a collection of randomly generated sparse environments. The *RiverSwim* environment consists of six states. The agent starts from the left side of the row and, in each state, can either swim left or right. Swimming to the right




Figure 1: *RiverSwim* Transition Model: the continuous (resp. dotted) arrows represent the transitions if action 1 (resp. 2) has been chosen.

(against the current of the river) is successful with probability 0.35; it leaves the agent in the same state with a high probability equal to 0.6, and leads him to the left with probability 0.05 (see Figure 1). On the contrary, swimming to the left (with the current) is always successful. The agent receives a small reward when he reaches the leftmost state, and a much larger reward when reaching the rightmost state – the other states offer no reward. This MDP requires efficient exploration procedures, since the agent, having no prior idea of the rewards, has to reach the right side to discover which is the most valuable state-action pair. The *SixArms* environment consists of seven states, one of which (state




Figure 2: *SixArms* Transition Model

0) is the initial state. From the initial state, the agent may choose one among six actions: the action $a \in \{1, \dots, 6\}$ leads to the state $x = a$ with probability p_a (see Figure 2) and let the agent in the initial state with probability $1 - p_a$. From all the other states, some actions deterministically lead the agent to the initial state while the others leave it in the current state. Staying in a state $x \in \{1, \dots, 6\}$, the agent receives a reward equal to R_x (see Figure 2), otherwise, no reward is received.

We compare the performance of the KL-UCRL algorithm to UCRL2 using 20 Monte-Carlo replications. For both algorithms, the constants C_P and C_R are settled to ensure that the upper bounds of the regret of Theorem 1 and Theorem 2 in [9] hold with probability 0.95. In the *SixArms* environment, the received rewards being deterministic, we slightly modify both algorithms so that the agent knows them beforehand. We observe in Figure 3 and 4 that the KL-UCRL algorithm accomplishes a smaller average regret than the UCRL2

algorithm in those benchmark environments. In both environments, it is crucial for the agent to learn that there is no action leading from some states to the most promising one: for example, in the *RiverSwim* environment, between one of the first four states and the sixth state.

In addition to those benchmark environments, a generator of sparse environments has been used to create 10-states and 5-actions environments with random rewards in $[0, 1]$. In these random environments, each state is connected with, on average, five other states (with transition probabilities drawn from a Dirichlet distribution). We reproduced the same experiments as in the previous environments and display the average regret in Figure 5.




Figure 3: Comparison of the regret of the UCRL2 and KL-UCRL algorithms in the RiverSwim environment.

6 Discussion

In this section, we expose the advantages of using a confidence ball based on the Kullback-Leibler divergence rather than an L^1 -ball, as proposed for instance in [9, 16], in the computation of the optimistic policy. This discussion aims at explaining and interpreting the difference of performance that can be observed in simulations. In KL-UCRL, optimism reduces to maximizing the linear function $V'q$ over a KL-ball (see (3)), whereas the other algorithms make use of an L^1 -




Figure 4: Comparison of the regret of the UCRL2 and KL-UCRL algorithms in the SixArms environment.




Figure 5: Comparison of the regret of the UCRL2 and KL-UCRL algorithms in randomly generated sparse environments.

ball:

$$\max_{q \in \mathbb{S}^{|X|}} V'q \quad \text{s.t.} \quad \|p - q\|_1 \leq \epsilon'. \quad (9)$$

Continuity

Consider an estimated transition probability vector p , and denote by q^{KL} (resp. q^1) the probability vector which maximizes Equation (3) (resp. Equation (9)). It is easily seen that q^{KL} and q^1 lie respectively on the border of the convex set $\{q \in \mathbb{S}^{|X|} : KL(p, q) \leq \epsilon\}$ and at one of the vertices of the polytope $\{q \in \mathbb{S}^{|X|} : \|p - q\|_1 \leq \epsilon'\}$. A first noteworthy difference between those neighborhoods is that, due to the smoothness of the KL-neighborhood, q^{KL} is continuous with respect to the vector V , which is not the case for q^1 .

To illustrate this, Figure 6 displays L^1 - and KL-balls around 3-dimensional probability vectors. The set of 3-dimensional probability vectors is represented by a triangle whose vertices are the vectors $(1, 0, 0)'$, $(0, 1, 0)'$ and $(0, 0, 1)'$, the probability vector p by a white star, and the vectors q^{KL} and q^1 by a white point. The arrow represents the direction of V 's projection on the simplex and indicates the gradient of the linear function to maximize. The maximizer q^1 can vary significantly for small changes of the value function, while q^{KL} varies continuously.

Unlikely transitions

Denote $i_m = \operatorname{argmin}_j V_j$ and $i_M = \operatorname{argmax}_j V_j$. As underlined by [9], $q_{i_m}^1 = \max(p_{i_m} - \epsilon'/2, 0)$ and $q_{i_M}^1 = \min(p_{i_M} + \epsilon'/2, 1)$. This has two consequences:

1. if p is such that $0 < p_{i_m} < \epsilon'/2$, then the vector $q_{i_m}^1 = 0$; so the optimistic model may assign a probability equal to zero to a transition that has actually been observed, which makes it hardly compatible with the optimism principle. Indeed, an optimistic MDP should not forbid transitions that really exists, even if they lead to states with small values;
2. if p is such that $p_{i_M} = 0$, then $q_{i_M}^1$ never equals 0; therefore, an optimistic algorithm that uses L^1 -balls will always assign positive probability to transitions to i_M even if this transition is impossible under the true MDP and if much evidence has been accumulated against the existence of such a transition. Thus, the exploration bonus of the optimistic procedure is wasted, whereas it could be used more efficiently to favor some other transitions.

This explains a large part of the experimental advantage of KL-UCRL observed in the simulations. Indeed, q^{KL} always assigns strictly positive probability to observed transitions, and eventually renounces unobserved transitions even if the target states have a potentially large value. Algorithm 2 works as follows: for all i such that $p_i \neq 0$, $q_i \neq 0$; for all i such that $p_i = 0$, $q_i = 0$ except if $p_{i_M} = 0$ and if $f(V_{i_M}) < \epsilon$, in which case $q_{i_M} = 1 - \exp(f(V_{i_M}) - \epsilon)$. But this is no longer the case when ϵ becomes small enough, that is, when sufficiently many observations are available. We illustrate those two important differences




Figure 6: The L^1 -neighborhood $\{q \in \mathbb{S}^3 : \|p - q\|_1 \leq 0.2\}$ (left) and KL-neighborhood $\{q \in \mathbb{S}^3 : KL(p, q) \leq 0.02\}$ (right) around the probability vector $p = (0.15, 0.2, 0.65)'$ (white star). The white points are the maximizers of equations (3) and (9) with $V = (0, 0.05, 1)'$ (up) and $V = (0, -0.05, 1)'$ (down).

in Figure 7, by representing the L^1 and KL neighborhoods together with the maximizers q^{KL} and q^1 , first if p_{i_m} is positive by very small, and second if p_{i_m} is equal to 0. Figure 8 also illustrates the latter case, by representing the evolution of the probability vector q that maximizes both (9) and (3) for an example with $p = (0.3, 0.7, 0)'$, $V = (1, 2, 3)'$ and ϵ decreasing from $1/2$ to $1/500$.

A Linear optimization over a KL-ball

This section explains how to solve the optimization problem of Equation (3). In [12], a similar problem arises in a different context, and a somewhat different solution is proposed for the case when the p_i are all positive. As a problem of maximizing a linear function under convex constraints, it is sufficient to consider the Lagrangian function

$$L(q, \lambda, \nu, \mu_1, \dots, \mu_N) = \sum_{i=1}^N q_i V_i - \lambda \left(\sum_{i=1}^N p_i \log \frac{p_i}{q_i} - \epsilon \right) - \nu \left(\sum_{i=1}^N q_i - 1 \right) + \sum_{i=1}^N \mu_i q_i .$$




Figure 7: The L^1 (left) and KL-neighborhoods (right) around the probability vector $p = (0, 0.4, 0.6)'$ (up) and $p = (0.05, 0.35, 0.6)'$ (down). The white point is the maximizer of the equations (3) and (9) with $V = (-1, -2, -5)'$ (up) and $V = (-1, 0.05, 0)'$ (down). We took, $\epsilon = 0.05$ (up), $\epsilon = 0.02$ (down) and $\epsilon' = \sqrt{2}\epsilon$.




Figure 8: Evolution of the probability vector q that maximizes both (3) (top) and (9) (bottom) with $p = (0.3, 0.7, 0)'$, $V = (1, 2, 3)'$ and ϵ decreasing from $1/2$ to $1/500$

If q is a maximizer, there exist $\lambda \in \mathbb{R}$, $\nu, \mu_i \geq 0$ ($i = 1 \dots N$) such that the following conditions are simultaneously satisfied:

$$\begin{cases} V_i + \lambda \frac{p_i}{q_i} - \nu + \mu_i = 0 \end{cases} \quad (10)$$

$$\begin{cases} \lambda \left(\sum_{i=1}^N p_i \log \frac{p_i}{q_i} - \epsilon \right) = 0 \end{cases} \quad (11)$$

$$\begin{cases} \nu \left(\sum_{i=1}^N q_i - 1 \right) = 0 \end{cases} \quad (12)$$

$$\begin{cases} \mu_i q_i = 0 \end{cases} \quad (13)$$

Let $Z = \{i, p_i = 0\}$. Conditions (10) to (13) imply that $\lambda \neq 0$ and $\nu \neq 0$. For $i \in \bar{Z}$, Equation (10) implies that $q_i = \lambda \frac{p_i}{\nu - \mu_i - V_i}$. Since $\lambda \neq 0$, $q_i > 0$ and then, according to (13), $\mu_i = 0$. Therefore,

$$\forall i \in \bar{Z}, \quad q_i = \lambda \frac{p_i}{\nu - V_i}. \quad (14)$$

Let $r = \sum_{i \in Z} q_i$. Summing on $i \in \bar{Z}$ and using equations (14) and (12), we have

$$\lambda \sum_{i \in \bar{Z}} \frac{p_i}{\nu - V_i} = \sum_{i \in \bar{Z}} q_i = 1 - r. \quad (15)$$

Using (14) and (15), we can write $\sum_{i \in \bar{Z}} p_i \log \frac{p_i}{q_i} = f(\nu) - \log(1 - r)$ where f is defined in (4). Then, q satisfies condition (11) if and only if $f(\nu) = \epsilon + \log(1 - r)$.

Consider now the case where $i \in Z$. Let $I^* = Z \cap \arg\max_i V_i$. Note that, for all $i \in Z \setminus I^*$, $q_i = 0$. Indeed, otherwise, μ_i should be zero, and then $\nu = V_i$ according to (10), which involves a possible negative denominator in (14). According to (13), for all $i \in I^*$, either $q_i = 0$ or $\mu_i = 0$. The second case implies that $\nu = V_i$ and $r > 0$ which requires that $f(\nu) < \epsilon$ so that (A) can be satisfied with $r > 0$. Therefore,

- if $f(V_i) < \epsilon$ for $i \in I^*$, then $\nu = V_i$ and the constant r can be computed solving equation $f(\nu) = \epsilon - \log(1 - r)$; the values of q_i for $i \in I^*$ may be chosen in any way such that $\sum_{i \in I^*} q_i = r$;
- if for all $i \in I^*$ $f(V_i) \geq \epsilon$, then $r = 0$, $q_i = 0$ for all $i \in Z$ and ν is the solution of the equation $f(\nu) = \epsilon$.

Once ν and r have been determined, the other components of q can be computed according to (14): we have that for $i \in \bar{Z}$, $q_i = \frac{(1-r)\tilde{q}_i}{\sum_{i \in \bar{Z}} \tilde{q}_i}$ where $\tilde{q}_i = \frac{p_i}{\nu - V_i}$.

B Properties of the f function

In this section, a few properties of function f defined in Equation (4) are stated, as this function plays a key role in the maximizing procedure of Section 3.2.

Proposition 2 f is a convex, decreasing mapping from $] \max_{i \in \bar{Z}} V_i; \infty[$ onto $]0; \infty[$.

Proof 1 Using Jensen's inequality, it is easily shown that the f function decreases from $+\infty$ to 0. The second derivative of f with respect to ν is equal to

$$-\sum_i \frac{p_i}{(\nu - V_i)^2} + \frac{2 \sum_i \frac{p_i}{(\nu - V_i)^3} \sum_i \frac{p_i}{\nu - V_i} - \left(\sum_i \frac{p_i}{(\nu - V_i)^2} \right)^2}{\left(\sum_i \frac{p_i}{\nu - V_i} \right)^2}.$$

If Z denotes a positive random value such that $\mathbb{P}\left(Z = \frac{1}{\nu - V_i}\right) = p_i$, then

$$f''(\nu) = \frac{2\mathbb{E}(Z^3)\mathbb{E}(Z) - \mathbb{E}(Z^2)\mathbb{E}(Z)^2 - \mathbb{E}(Z^2)^2}{\mathbb{E}(Z)^2}.$$

Using Cauchy-Schwartz inequality, we have $\mathbb{E}(Z^2)^2 = \mathbb{E}(Z^{3/2}Z^{1/2})^2 \leq \mathbb{E}(Z^3)\mathbb{E}(Z)$. In addition $\mathbb{E}(Z^2)^2 \geq \mathbb{E}(Z^2)\mathbb{E}(Z)^2$. These two inequalities show that $f''(\nu) \geq 0$.

As mentioned in Section 3.2, Newton's method can be applied to solve the equation $f(\nu) = \epsilon$ for a fixed value of ϵ . When ϵ is close to 0, the solution of this equation is quite large and an appropriate initialization accelerates convergence. Using a second-order Taylor's-series approximation of the function f , it can be seen that, for ν near ∞ , $f(\nu) = \frac{\sigma_{p,V}}{2\nu^2} + o(\frac{1}{\nu^2})$, where $\sigma_{p,V} = \sum_i p_i V_i^2 - (\sum_i p_i V_i)^2$. The Newton iterations can thus be initialized by taking $\nu_0 = \sqrt{\sigma_{p,V}/(2\epsilon)}$.

References

- [1] P. Auer, T. Jaksch, and R. Ortner. Near-optimal regret bounds for reinforcement learning. *Advances in Neural Information Processing Systems*, 21, 2009.
- [2] P. Auer and R. Ortner. Logarithmic online regret bounds for undiscounted reinforcement learning. *Advances in Neural Information Processing Systems*, page 49, 2007.
- [3] P. Bartlett and A. Tewari. REGAL: A Regularization based Algorithm for Reinforcement Learning in Weakly Communicating MDPs. *Annual Conference on Uncertainty in Artificial Intelligence*, 2009.
- [4] R. Brafman and M. Tennenholtz. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3:213–231, 2003.
- [5] A. Burnetas and M. Katehakis. Optimal adaptive policies for Markov decision processes. *Mathematics of Operations Research*, pages 222–255, 1997.

- [6] E. Even-Dar, S. Mannor, and Y. Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.
- [7] A. Garivier and F. Leonardi. Context tree selection: A unifying view. preprint, 2010.
- [8] A. Garivier and E. Moulines. On upper-confidence bound policies for non-stationary bandit problems. *Arxiv preprint arXiv:0805.3415*, 2008.
- [9] T. Jaksch, R. Ortner, and P. Auer. Near-optimal Regret Bounds for Reinforcement Learning. *Journal of Machine Learning Research*, 11:1563–1600, 2010.
- [10] M. Kearns and S. Singh. Near-optimal reinforcement learning in polynomial time. *Journal of machine learning*, 49(2-3):209–232, 2002.
- [11] T. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [12] A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- [13] M. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc. New York, NY, USA, 1994.
- [14] A. Strehl and M. Littman. An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.
- [15] R. Sutton and A. Barto. *Reinforcement learning: An introduction*. The MIT press, 1998.
- [16] A. Tewari and P. Bartlett. Optimistic linear programming gives logarithmic regret for irreducible MDPs. *Advances in Neural Information Processing Systems*, 20:1505–1512, 2008.

